

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

SUPERSPEED, L.L.C.,	§
Plaintiff,	§
	§
	§
v.	§ CIVIL ACTION NO. 2-07-CV-89
	§
INTERNATIONAL BUSINESS MACHINES	§
CORPORATION,	§
Defendant.	§
	§
	§

MEMORANDUM OPINION AND ORDER

After considering the submissions and the arguments of counsel, the Court issues the following order concerning the claim construction issues:

I. Introduction

Plaintiff SuperSpeed, L.L.C. (“SuperSpeed”) alleges Defendant International Business Machines Corporation (“IBM”) infringes five of its patents: U.S. Patent Nos. 5,577,226 (“the ‘226 Patent”), 5,918,244 (“the ‘244 Patent”), 7,017,013 (“the ‘013 Patent”), 7,039,767 (“the ‘767 Patent”), and 7,111,129 (“the ‘129 Patent”). Each of the five SuperSpeed Patents claims priority to a common parent application: U.S. Application No. 08/238,815. The application was filed May 6, 1994, and issued as the ‘226 Patent on November 19, 1996. Each of the other Patents stem from continuation applications claiming the same priority date.¹ Both SuperSpeed and IBM agree that the terms should be construed the same across all asserted patents.

Of the five patents SuperSpeed alleges IBM infringes, the ‘226 and ‘244 Patents were asserted against Oracle Corporation in a suit filed in the Southern District of Texas (“the Oracle

¹The patents share a common specification, therefore, references to the specification will cite the ‘226 patent.

litigation”). Judge Gilmore presided over the Markman Hearing in that case and has issued a claim Construction Order with respect to certain terms in the ‘226 and ‘244 Patents. Several of the disputed terms in this case, as discussed below, were construed in the Oracle litigation.²

II. Background of the Technology

The five SuperSpeed patents disclose a system that effectuates faster access to shared data on a network. The system also provides for improving “caching coherency.” “Cache coherency” is the term used to describe maintaining the integrity of shared data on a network when that data has been cached on multiple computers on the network. The preferred embodiment describes a computer on a network that communicates with shared I/O devices on the network (e.g., disks or other storage devices). Conventionally, to speed up the retrieval of data from a storage device, the data used by the computer may be temporarily loaded or “cached” in memory that is local to the computer. This local memory is typically dynamic random access memory (“RAM”), also known as “system memory,” that is internal to the computer itself.

SuperSpeed’s patents describe the prior art as being made up of network caching systems that stored data using a “one-size-fits-all” approach. SuperSpeed’s inventions change that system by creating one that uses multiple “bucket sizes” for storing data. The network caching software of the SuperSpeed Patents allows for the use of multiple data bucket sizes, each with a different amount of memory space. These data buckets can be customized to accept small, medium and large pieces of data, as opposed to a single size. This innovation maximizes system resources by accommodating large pieces of data (in the large bucket size) and, at the same time, avoiding excessive waste when

²That case settled before judgment and the parties are not arguing that either side is bound by Judge Gilmore’s constructions. Her constructions are, however, instructive in the instant case.

smaller pieces of data are cached.

In addition, SuperSpeed's invention discloses a way to speed up data access by changing the way that cache coherency is maintained. Prior to SuperSpeed's invention, if a computer changed data that was accessible to multiple computers on the network, that computer must broadcast a message on the network that the data had been changed. This preserved the integrity of that data on other computers that may have cached a copy of it. SuperSpeed's inventions minimize network communication traffic by keeping track of the computers that may cache data from a particular storage device, and then notifying only those computers when a data block from that storage device has been modified.

III. General Principles Governing Claim Construction

“A claim in a patent provides the metes and bounds of the right which the patent confers on the patentee to exclude others from making, using or selling the protected invention.” *Burke, Inc. v. Bruno Indep. Living Aids, Inc.*, 183 F.3d 1334, 1340 (Fed. Cir. 1999). Claim construction is an issue of law for the court to decide. *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 970-71 (Fed. Cir. 1995) (en banc), *aff’d*, 517 U.S. 370 (1996).

To ascertain the meaning of claims, the court looks to three primary sources: the claims, the specification, and the prosecution history. *Markman*, 52 F.3d at 979. Under the patent law, the specification must contain a written description of the invention that enables one of ordinary skill in the art to make and use the invention. A patent's claims must be read in view of the specification, of which they are a part. *Id.* For claim construction purposes, the description may act as a sort of dictionary, which explains the invention and may define terms used in the claims. *Id.* “One purpose for examining the specification is to determine if the patentee has limited the scope of the claims.”

Watts v. XL Sys., Inc., 232 F.3d 877, 882 (Fed. Cir. 2000).

Nonetheless, it is the function of the claims, not the specification, to set forth the limits of the patentee's claims. Otherwise, there would be no need for claims. *SRI Int'l v. Matsushita Elec. Corp.*, 775 F.2d 1107, 1121 (Fed. Cir. 1985) (en banc). The patentee is free to be his own lexicographer, but any special definition given to a word must be clearly set forth in the specification. *Intellicall, Inc. v. Phonometrics*, 952 F.2d 1384, 1388 (Fed. Cir. 1992). And, although the specification may indicate that certain embodiments are preferred, particular embodiments appearing in the specification will not be read into the claims when the claim language is broader than the embodiments. *Electro Med. Sys., S.A. v. Cooper Life Sciences, Inc.*, 34 F.3d 1048, 1054 (Fed. Cir. 1994).

This court's claim construction decision must be informed by the Federal Circuit's decision in *Phillips v. AWH Corporation*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). In *Phillips*, the court set forth several guideposts that courts should follow when construing claims. In particular, the court reiterated that "the *claims* of a patent define the invention to which the patentee is entitled the right to exclude." 415 F.3d at 1312 (emphasis added) (*quoting Innova/Pure Water, Inc. v. Safari Water Filtration Systems, Inc.*, 381 F.3d 1111, 1115 (Fed. Cir. 2004)). To that end, the words used in a claim are generally given their ordinary and customary meaning. *Id.* The ordinary and customary meaning of a claim term "is the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention, i.e., as of the effective filing date of the patent application." *Id.* at 1313. This principle of patent law flows naturally from the recognition that inventors are usually persons who are skilled in the field of the invention. The patent is addressed to and intended to be read by others skilled in the particular art. *Id.*

The primacy of claim terms notwithstanding, *Phillips* made clear that “the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification.” *Id.* Although the claims themselves may provide guidance as to the meaning of particular terms, those terms are part of “a fully integrated written instrument.” *Id.* at 1315 (*quoting Markman*, 52 F.3d at 978). Thus, the *Phillips* court emphasized the specification as being the primary basis for construing the claims. *Id.* at 1314-17. As the Supreme Court stated long ago, “in case of doubt or ambiguity it is proper in all cases to refer back to the descriptive portions of the specification to aid in solving the doubt or in ascertaining the true intent and meaning of the language employed in the claims.” *Bates v. Coe*, 98 U.S. 31, 38 (1878). In addressing the role of the specification, the *Phillips* court quoted with approval its earlier observations from *Renishaw PLC v. Marposs Societa’ per Azioni*, 158 F.3d 1243, 1250 (Fed. Cir. 1998):

Ultimately, the interpretation to be given a term can only be determined and confirmed with a full understanding of what the inventors actually invented and intended to envelop with the claim. The construction that stays true to the claim language and most naturally aligns with the patent’s description of the invention will be, in the end, the correct construction.

Consequently, *Phillips* emphasized the important role the specification plays in the claim construction process.

The prosecution history also continues to play an important role in claim interpretation. The prosecution history helps to demonstrate how the inventor and the PTO understood the patent. *Phillips*, 415 F.3d at 1317. Because the file history, however, “represents an ongoing negotiation between the PTO and the applicant,” it may lack the clarity of the specification and thus be less useful in claim construction proceedings. *Id.* Nevertheless, the prosecution history is intrinsic

evidence. That evidence is relevant to the determination of how the inventor understood the invention and whether the inventor limited the invention during prosecution by narrowing the scope of the claims.

Phillips rejected any claim construction approach that sacrificed the intrinsic record in favor of extrinsic evidence, such as dictionary definitions or expert testimony. The *en banc* court condemned the suggestion made by *Texas Digital Systems, Inc. v. Telegenix, Inc.*, 308 F.3d 1193 (Fed. Cir. 2002), that a court should discern the ordinary meaning of the claim terms (through dictionaries or otherwise) before resorting to the specification for certain limited purposes. *Id.* at 1319-24. The approach suggested by *Texas Digital*—the assignment of a limited role to the specification—was rejected as inconsistent with decisions holding the specification to be the best guide to the meaning of a disputed term. *Id.* at 1320-21. According to *Phillips*, reliance on dictionary definitions at the expense of the specification had the effect of “focus[ing] the inquiry on the abstract meaning of words rather than on the meaning of the claim terms within the context of the patent.” *Id.* at 1321. *Phillips* emphasized that the patent system is based on the proposition that the claims cover only the invented subject matter. *Id.* What is described in the claims flows from the statutory requirement imposed on the patentee to describe and particularly claim what he or she has invented. *Id.* The definitions found in dictionaries, however, often flow from the editors’ objective of assembling all of the possible definitions for a word. *Id.* at 1321-22.

Phillips does not preclude all uses of dictionaries in claim construction proceedings. Instead, the court assigned dictionaries a role subordinate to the intrinsic record. In doing so, the court emphasized that claim construction issues are not resolved by any magic formula. The court did not impose any particular sequence of steps for a court to follow when it considers disputed claim

language. *Id.* at 1323-25. Rather, *Phillips* held that a court must attach the appropriate weight to the intrinsic sources offered in support of a proposed claim construction, bearing in mind the general rule that the claims measure the scope of the patent grant. The court now turns to a discussion of the disputed claim terms.

IV. Terms in Dispute

A. Agreed Constructions

The parties have stipulated to the construction of the following terms in the claims:

“Bucket” means “a fixed-sized area in the cache where disk data is stored.”

“Remote nodes” means “other computers on the network.”

“Cache software,” “cache driver,” and “cache program” mean “a software program that creates or controls a cache.”

“System memory” means “the system main random access memory of a computer.”

“Write instruction” means “an operation that initiates a transfer of data to an I/O device.”

“Cache” as used as a noun means “a portion of system main memory used for temporary storage of I/O data in a cache,” and as used as a verb means “to store I/O data in a cache.”

“I/O device” means “disk or other persistent storage device.”

B. Disputed Constructions

1. “Suitable bucket size”

SuperSpeed proposes the term should be construed as “bucket size selected based upon the size of the data to be cached.” IBM’s proposed construction is “bucket size selected based upon the size of the I/O access.” The dispute concerns how the system and method determine which bucket size is “suitable.”

IBM argues that the size is determined based on the size of the data transfer from the I/O device. SuperSpeed argues that “disk access” is ambiguous. Further, SuperSpeed argues, under IBM’s construction, the patent would not achieve the invention’s goal of using the cache memory effectively. IBM supports its position by citing to several parts of the specification where the preferred embodiment chooses where the data is put based upon “disk access”: “disk data is . . . copied into an available cache bucket depending upon size fit.” [‘226 Patent, 2:24-55], “total cache bucket size associated with it for small, medium, and large, disk access sizes.” [‘226 Patent, 1:54-57], “[u]sing the size fo the read data access the cache driver selects which of the three caches . . . the data transfer fits.” [‘226 Patent, 4:10-12].

Despite the illustrative passages, the invention’s goal is to minimize wasted space and speed up retrieval of cached data. There is no support in the intrinsic evidence to support the limitation imposed by IBM’s construction. Assuming *arguendo* that the disclosed embodiment does make some decision based on a “disk access,” the invention is not limited to the preferred embodiment. The inventor intended the “suitable bucket size” to be a bucket size that was the most appropriate size available with the least amount of wasted space. The term “suitable bucket size,” therefore means “a bucket size selected based on the size of the data to be cached.”

2. “Selecting one of the plurality of cache data bucket sizes available in one of the computers”

The dispute here is whether the court should impose a sequential limitation on method claim 5. SuperSpeed maintains this term does not need construction, while IBM proposes “this step in the method must follow ‘receiving, in one of the computers, a write instruction having data to be written into addresses on one of said I/O devices.’” Claim 5 is reproduced below:

A method for accelerating access to data on a network comprising:

- providing a plurality of computers on the network, each with cache software for creating caches with a plurality of cache data bucket sizes;
- receiving, in one of the computers, a write instruction having data to be written into addresses on one of said I/O devices;
- selecting one of the plurality of cache data bucket sizes available in one of the computers;
- writing the data into a data bucket of the selected data bucket size responsive to the write instruction; and
- communicating over the network with remote caches to invalidate cache data corresponding to any of the addresses.

[‘013 Patent, 26:35-49]

The Federal Circuit has made clear that “[u]nless the steps of a method actually recite an order, the steps are not ordinarily construed to require one.” *Interactive Gift Exp., Inc. v. Compuserve Inc.*, 256 F.3d 1353, 1342. “*Interactive Gift* recites a two-part test for determining if the steps of a method claim that do not otherwise recite an order, must nonetheless be performed in the order in which they are written.” *Altiris Inc. v. Symantec Corp.*, 318 F.3d 1363, 1370 (Fed. Cir. 2003). First, courts must “look to the claim language to determine if, as a matter of logic or grammar, they must be performed in the order written . . . If not, [courts] next look to the rest of the specification to determine whether it ‘directly or implicitly requires such a narrow construction.’” *Id.* (internal citations omitted).

IBM argues that logically, “without first receiving a write instruction identifying the size of the data to be written to the I/O device, the bucket size that fits the data cannot be selected.” IBM points to two places in the specification to support its argument. First, “[u]sing the size of the read data access the cache driver selects which of the three caches, small, medium or large, the data transfer fits.” [‘226 Patent, 4:10-12] Second, “[t]he program matches the byte count size of the

intercepted read I/O data transfer against the three cache sizes, small, medium, or large, attempting to choose which of the three TCH cache control structures this read I/O data will be targeted at.” [‘226 Patent, 18-28-32]

IBM’s arguments merely relate to the preferred embodiment. The fact that the steps happen in a specific sequence in the preferred embodiment is not enough to impose that limitation on the claim. *Altiris*, 318 F.3d at 1371. IBM has not made a persuasive argument why the cache program must always determine and select a suitable bucket size after receiving the write instruction. IBM’s argument fails both prongs of the *Interactive Gift* test. The grammar or logic of the claim language itself does not require such a limitation. Additionally, the patent as a whole does not directly or implicitly require such a narrow construction. The court, therefore, rejects IBM’s suggested limitation.

3. “**Invalidate**”

“Invalidate data” was construed by Judge Gilmore in the Oracle Litigation. Judge Gilmore construed “invalidate data” to mean “to indicate previously cached data has been modified.” IBM maintains this is the proper construction. SuperSpeed, however, would have the court construe the term to mean “to indicate that a portion of data in a cache is no longer up to date.” Here, the parties choose to have only “invalidate” construed because it is used with terms other than “data.” For example, claim 4 of the ‘013 patent uses “communicating invalidate messages.”

SuperSpeed argues that Judge Gilmore’s construction of “invalidate data” addressed different issues, such that her construction provides no real guidance for this court in construing “invalidate” alone. SuperSpeed raises two issues with IBM’s construction. First, SuperSpeed argues IBM’s construction requires that the copy of the data in the cache be modified, which excludes the preferred

embodiment. Second, SuperSpeed argues that IBM’s construction impermissibly limits the claim by requiring “invalidate” to only encompass data that has been modified, and not data that is being modified. SuperSpeed’s first argument, that IBM’s construction requires that some physical copy of the data be modified, is rejected. IBM’s construction does no such thing.

SuperSpeed’s second argument concerns how the parties understand “has been modified.” SuperSpeed maintains that the construction could be read to require that the modification be complete, which is an unnecessary limitation. SuperSpeed suggests the invalidation could occur before, or simultaneous with, any modification; the specification does not require that the invalidation occur after the write operation to the I/O device has been completed. The specification provides for cached data to be invalidated as part of the process of writing to disk: “[i]n accordance with the embodiment of the invention, when a write access is performed to a disk which is being cached and the disk data area being written was previously read into the cache, i.e. an update operation on the disk data, the current cache buckets for the previous read disk data area are invalidated on all computers on the network. [‘226 Patent, 2:45-50]. Nothing in the specification precludes invalidation before the modification has been fully completed. For example, the invalidate message could be sent simultaneously with the write command, or immediately preceding it.

Additionally, the term should not be construed to require a temporal limitation, because some of the claims with the term themselves import the limitation. *See Phillips v. AWH Corp.*, 415 F.3d at 1314 (holding claim context is highly instructive when construing claim terms). For example, claim 4 of the ‘013 patent requires that the invalidation occur after the data is written to the I/O device: “after data is written to the shared I/O device, communicating invalidate messages to only the computers in the list . . .” Claim 20 of the ‘013 patent has a similar temporal requirement. If the

inventor would have intended the term to import a sequential limitation, there would have been no need to do so in those claims separately.

Judge Gilmore did not have this issue before her when construing “invalidate data” in the Oracle Litigation. Before Judge Gilmore were competing definitions, one of which was SuperSpeed’s current proposition, which she rejected. Oracle’s, which she also rejected, was “make data unusable.” Judge Gilmore crafted her own construction in light of the specification. Judge Gilmore found that “the summary of the invention and specifications are consistent with a claim construction of ‘invalidate data’ that indicates a change or modification of previous data with more recent data.” *Oracle Litigation* (citing ‘226 Patent, at 2:13-20, 2:45-50; ‘244 Patent, at 2:15-22, 2:46-52).

The court agrees with SuperSpeed that “invalidate,” standing alone, does not import a sequential limitation. The court, therefore, construes “invalidate” to mean “to indicate the modification of previously cached data.”

4. “Network”

IBM proposes the term means “a group of computers and peripheral devices connected together so that they can communicate with each other.” SuperSpeed proposes the term means “communication facilities that link points at which computers or devices may be connected.” SuperSpeed’s construction was the agreed construction in the Oracle Litigation, and was adopted by Judge Gilmore in her Claim Construction Ruling. There is nothing in the intrinsic record that supports IBM’s proposed construction. As SuperSpeed points out, computers or peripheral devices may be “connected to” or put “on” a network. *See, e.g.*, ‘226 claim 27 (“a computer connected to the network . . .”; ‘013 claim 4 (“providing a plurality of computers on the network . . .”). In the context of

this patent, however, the devices or peripherals themselves do not comprise the network.

IBM heavily relies on claim 1 of the ‘767 patent: “a computer network that comprises: one or more I/O devices on which data may be stored in files; and multiple computers coupled together.” That reliance is misplaced, however, because “computer network” is different from “network” as used alone in the other claims.³

There is no support for deviating from Judge Gilmore’s construction. In the context of this patent, the court construes “network” to mean “communication facilities that link points at which computers or devices may be connected.”

5. **“intercepting”**

SuperSpeed proposes the term means “to stop, deflect, or interrupt the progress or intended course of.” IBM proposes the term means “cutting off from the intended destination.” In the Oracle Litigation, the parties agreed to SuperSpeed’s current proposal, and Judge Gilmore adopted that agreement in her claim construction ruling. IBM’s construction would necessarily mean that the instruction does not reach its final destination. This is not what is disclosed in the patent.

The read and write instructions must be intercepted for the cache driver to determine which data should be read from or written to the cache. The read instruction, for example, must be sent to the disk when the data is not found in the cache. [‘226 Patent, 4:47-51]. The write instruction must be sent to the I/O device to write the data to the non-volatile storage. As SuperSpeed points out, the “interception” does not cut off the instruction from the target, but only stops it or interrupts it before

³*Id. see also, e.g.,* “A caching system comprising: a network; . . .” [‘244 Patent, cl. 15, 25]; A caching system comprising: a network; a first computer . . .” [‘244 Patent, cl. 25]

sending it to the appropriate I/O device. [‘226 Patent, 4:47-51] (noting that, after interception, when the data sought by a read instruction is not in the cache “the disk is accessed normally for the reqd data”); and ‘226 23:32-45 (noting that, after interception, a write instruction is ultimately “sent to the disk I/O device”). SuperSpeed argues, therefore, IBM’s proposed construction excludes the preferred embodiment. When a construction excludes the preferred embodiment, it “is rarely, if ever, correct.”

Vitronics, 90 F.3d at 1583.

IBM responds its proposed construction does not exclude the preferred embodiment. It argues that the “instruction” is intercepted, and then a new instruction is sent after the cache driver determines what to do with the data. IBM’s argument is unpersuasive. There is no support for IBM’S proposition that the original read or write instruction can not be passed along by the cache driver. The above passages suggest as much. IBM is correct that the specification describes a situation where the driver modifies the instruction. *See* ‘226 Fig. 5G at block 496 (showing that the invention will “adjust [the] I/O transfer request to intercept completion at ‘READ COMPLETE’” for a read instruction for which no data is in the cache). That, however, is not necessary in all cases.

The term “intercepting,” therefore, means “to stop, deflect, or interrupt the progress or intended course of.”

6. “computer using an n-bit architecture” / “Computer using a m-bit architecture”

SuperSpeed proposes the term means “a computer whose general purpose registers are n-bits wide.”⁴ IBM proposes the term means “a computer that uses an address that has a length of ‘n’ bits

⁴ The same analysis applies to “m-bit architecture.” The court will refer to “n-bit architecture” for simplicity’s sake.

to address data stored in main memory.” SuperSpeed cites to the intrinsic record to support its proposal: “[t]he presently preferred embodiment of the invention uses remote disks that are connected by the Open VMS VMScluster and VAXcluster software. The VMScluster software is operable on 64-bit or 32-bit architecture computer systems. The VAXcluster software only permits 32-bit computers.” [‘244 Patent, 3:26-30]

SuperSpeed has presented extrinsic evidence supporting that the “VAXcluster software,” not surprisingly, was typically used on VAX computers that had general-purpose registers that were 32-bits wide. These computers, however, used an address *with only 28 bits* to address data stored in main memory. Bennett Dec., ¶ 43. The Alpha AXP computer referred to in the specification, which had a “64-bit architecture,” used an address *with only 48 bits* to address data stored in main memory. Bennett Dec., ¶ 44. SuperSpeed points to other relevant extrinsic evidence to prove that it was well understood in the art that a computer having a 32 or 64 bit architecture meant the computer had general purpose registers that were 32 or 64 bits wide respectively. *See* Bennett Dec., ¶ 42; *see also* Hennessey & Patterson’s *Computer Architecture: A Quantitative Approach*. This is the bit-length through which a computer manipulates data in its CPU. For example, Hennessey and Patterson describe the IBM System 360 as a “32-bit machine” even though they note that it “had a 24-bit address space.” Hennessey & Patterson at 148.

The intrinsic and extrinsic record support that a computer with an “n-bit architecture” or a “m-bit architecture” refers to a computer’s general purpose register through which it manipulates data. The terms “computer using an n-bit architecture” and “computer using an m-bit architecture,” therefore, mean “a computer whose general purpose registers are n-bits wide.” and “a computer whose general purpose registers are m-bits wide,” respectively.

7. **“Means for intercepting a read instruction to one of said plurality of I/O devices from the compute on which said cache driver resides” and “means for intercepting a write instruction to one of said plurality of I/O devices”**

Both parties agree these are mean-plus-function claim limitations and that the recited functions should be given the meaning “intercepting a read instruction to one of said plurality of I/O devices from the computer on which said cache driver resides,” and “intercepting a write instruction to one of said plurality of I/O devices,” respectively. The parties disagree about the corresponding structure, however. SuperSpeed proposes the corresponding structure should be “a computer executing software that replaces the I/O entry point for an I/O device with the entry point to the cache driver.” IBM proposes “the i/o intercept global” program flow, steps and data structures disclosed or referred to by Fig. 2B and Col. 6, Ins. 35-64 of the ‘226 patent.”

SuperSpeed argues that IBM’s proposed structure is too broad; the “i/o intercept global” subroutine wrongly includes structure that is not clearly linked to the claimed invention. For example, portions of the “i/o intercept global” subroutine relate to determining whether an I/O device is supported, building and modifying TCB disk control structures, and clearing the remote system access list within the TCB. [‘226 Patent, 6:27-30, 6:40-49; Fig. 2B boxes 76-80]. SuperSpeed argues that IBM’s construction is too narrow at the same time. It argues IBM’s construction ignores the “i/o intercept device” subroutine, which relates to subsequent interceptions once the cache driver is operational. [‘226 Patent, 16:1-2; 16-40-44; Fig. 5A box 404].

IBM argues that SuperSpeed’s construction is improper under *Aristocrat Technologies Australia PTY Ltd. v. Int’l Game Tech.*, 521 F.3d at 1328, 1334 (Fed. Cir. 2008), because it simply describes the outcome of the interception algorithm—replacing the I/O entry point rather than

identifying the algorithm itself. In *Aristocrat*, the court found that language *in the claim itself* was not a sufficient algorithmic structure for the disputed term “game control means.”⁵ *Id.* The court found that the language Aristocrat argued was the equation, actually only described “the result of practicing the . . . function. That is, the equation is not an algorithm that describes how the function is performed, but is merely a mathematical expression that describes the outcome of performing the function.” *Id.*

SuperSpeed responds that the structure may be construed as a simple description of the detailed algorithm in that patent. *Harris Corp. v. Ericsson*, 417 F.3d 1241, 1254 (Fed. Cir. 2005) (construing the structure for “time domain processing means” as being “a microprocessor programmed to carry out a two-step algorithm in which the processor calculates generally nondiscrete estimates and then selects the discrete value closest to each estimate”). Therefore, SuperSpeed argues, its summary of the part of the algorithm that “intercepts” in the patent suffices as recited means corresponding to the function. This is persuasive. The *Harris Corp.* court found “symbol processor” was too broad of a structure to correspond to the recited function in the district court’s claim construction. *Id.* at 1253-55 (citing *WMS Gaming, Inc. v. Int’l Game Tech.*, 184 F.3d 1339, 1348-49 (Fed. Cir. 1999)) (explaining “WMS Gaming restricts computer-implemented means-plus-function terms to the algorithm disclosed in the specification.”). The *Harris* court, however, did not mandate that the entire detailed algorithm be included as part of the recited means of the means plus function term at issue. *See id.* at 1254 (analyzing the disclosed algorithm to arrive at a construction that would inform the jury on whether the claims of the patent in issue “cover systems that implement

⁵ The language in question was “defining a set of predetermined arrangements for a current game comprising each possible combination of the symbol position selected by the player which have one and only one symbol position in each column of the display means.”

either a one-step or two-step process"). Here too, an analysis of the algorithm detailed in the '226 patent informs the court that the construction proposed by the plaintiff is an accurate summary of the relevant parts of the disclosed algorithms. The Summary of the Invention provides:

When an I/O device is found to be one of the disk device types supported by the cache software of the invention (70), **the program intercepts the I/O entry point for the I/O device (74) by replacing it with an entry into the program routine "process io" (400, FIG. 5A) within the cache software of the invention.** A TCB (16, FIG. 1B) disk control structure for the disk I/O device is built . . .

SuperSpeed's construction provides sufficient detail to satisfy the requirements of the holding in *Harris Corp.*, 417 F.3d at 1254. Unlike the district court's construction in *Aristocrat*, "replacing the I/O entry point with an entry point into a different routine" is not the result of the algorithm, it is how the algorithm *performs* the function of intercepting the instruction—"stop[ping], deflect[ing], or interrupt[ing] the progress or intended course of" the instruction. IBM also argues that SuperSpeed's construction broadens the claim by allowing the "means for intercepting . . ." to be in any software and not just the cache driver which is required by the patent claims. The claim language is "[t]he caching system of claim 1 wherein *each of said cache drivers further includes* means for intercepting . . ." (emphasis added) The plain language of the claim requires the cache driver. Therefore, to include it in the construction would be redundant.

The corresponding structure for the two terms, therefore, is be: "a computer executing a software algorithm that replaces the I/O entry point for an I/O device with the entry point to the cache driver."⁶

⁶Pursuant to 35 U.S.C. § 112 ¶ 6, the literal scope of the claim includes such structure and equivalents thereof.

8. “means for . . . reading data from the cache when the read instruction relates to addresses corresponding to data in the cache”

Both parties agree this is a means-plus-function claim limitation and that the recited function should be given the meaning “reading data from the cache when the read instruction relates to addresses corresponding to data in the cache.” The parties disagree about the corresponding structure, however. SuperSpeed proposes the corresponding structure should be “a computer executing software that copies the data requested by a read instruction from the corresponding cache data bucket into the system memory location specified in the read instruction. IBM proposes “the ‘read data’ and ‘read cache hit’ program flow, steps, and data structures disclosed or referred to by Figs. 5c, 5e, and 5i, col. 17, l. 43-col.19, l. 6, and col. 20, l.60-col. 21, l. 15 of the ‘226 patent.

SuperSpeed argues that IBM’s proposed construction improperly includes structure for determining whether a “cache hit” has occurred (i.e., the “read data” subroutine). The “read data” subroutine is used to determine whether or not the requested data is in the cache, and performs “initial checks over the disk I/O device and its intercepted read I/O data transfer,” among other things. ‘226 17:43-19:14; Figs 5C-5E. Further, SuperSpeed argues that IBM’s proposed construction includes the “read cache hit” subroutine, which is contended to be unrelated to the function of reading data from the cache. SuperSpeed argues its proposed structure is correct because it claims only the structure necessary to read data from the cache after other structure determines whether the data is in the cache.

IBM argues that both determining whether there is a cache hit and reading data from a cache is necessary to perform the entire recited function. IBM further argues that SuperSpeed’s proposed corresponding structure does not identify any specific algorithm, but only a slightly more detailed account of the software functionality. *See Aristocrat*, 521 F.3d 1334. IBM’s first argument is

persuasive. The claim does not include any other means or structure for determining whether there was actually a cache hit. To give meaning to the functional language in the claim, determining whether there is a cache hit must be part of the corresponding structure.

SuperSpeed also argues including the “read cache hit” subroutine is improper because it is unrelated to the function of reading data from the cache. For example, the “read cache hit” subroutine describes structure for: (1) moving a TCMH cache memory block to the front of the least recently used queue, (2) checking whether the cached data is valid, (3) incrementing the “cache hit count” and (4) sending an I/O completion signal when the data is read successfully. ‘226 20:61-63, 20:66-21:5; Fig 5I, boxes 548, 552-556. SuperSpeed’s expert opines that while that structure may be necessary to enable the claimed invention to work, it should not be identified as the “corresponding structure” because it is not necessary to perform the recited function. Bennett Dec., ¶ 53

Although SuperSpeed’s argument is not without some force, the patentee in this instance claimed narrowly, invoking the dictates of § 112 ¶6. To perform the entire function, the structure for determining whether a cache hit or cache miss is necessary. Although IBM’s proposal includes some structure necessary for the invention to work, all of that structure is also disclosed as linked to performing the recited function, as set forth in the claim and construed by the court. The corresponding structure for the recited function, therefore, is “the ‘read data’ and ‘read cache hit’ program flow, steps, and data structures disclosed or referred to by Figs. 5c, 5e, and 5i, col. 17, l. 43- col.19, l. 6, and col. 20, 1.60-col. 21, l. 15 of the ‘226 patent.”⁷

⁷There may be better or more efficient structure to perform the recited function, but this is the structure disclosed in the specification. The court’s construction does not necessarily preclude those better or more efficient structures, however, provided they are equivalents. 35 U.S.C. § 112 ¶ 6.

9. “means for writing data into the cache when the read instruction relates to addresses that do not correspond to any data in the cache”

This term is closely related to the previous term. Here, however, instead of reading data into the cache when there is a cache hit, the term is concerned with writing data into the cache when there is a cache miss. Both parties agree this is a mean-plus-function claim limitation and that the recited function should be given the meaning “writing data into the cache when the read instruction relates to addresses that do not correspond to any data in the cache.” The parties disagree, however, on the corresponding structure. SuperSpeed proposes the corresponding structure should be “a computer executing software that copies the data requested by a read instruction from a location in system memory into a cache data bucket.” IBM proposes, “the ‘read data’ and ‘read miss’ program flow, steps, and data structures disclosed or referred to by Figs. 5c, 5e, and 5f-5h, col. 17, l. 43-col. 19, l. 6, and col. 19, l. 14-col. 20, l.54 of the ‘226 patent.”

The parties raise the same arguments for this term as they did for the previous term. For the same reasons discussed above, the claimed structure should include not only the function for writing the data into the cache, but also for determining if there is a “cache hit” or “cache miss.” SuperSpeed again argues there are certain “housekeeping” aspects included in the structure proposed by IBM. Those aspects, however, are all disclosed as part of the routines linked to performing the recited function in the manner construed by the court. The corresponding structure for the recited function, therefore, is “the ‘read data’ and ‘read miss’ program flow, steps, and data structures disclosed or referred to by Figs. 5c, 5e, and 5f-5h, col. 17, l. 43-col. 19, l. 6, and col. 19, l. 14-col. 20, l.54 of the ‘226 patent, and equivalents thereof.”

10. “means for using the computer communication channels to invalidate data in the caches of any computer that is caching said one of said plurality of I/O devices”

Both parties agree this is a mean-plus-function claim limitation and that the recited function should be given the meaning “using the computer communication channels to invalidate data in the caches of any computer that is caching said one of said plurality of I/O devices.” The parties disagree, however, on the corresponding structure. SuperSpeed proposes that the corresponding structure should be “a computer executing software that sends a message indicating the invalidity of certain data found on an I/O device to the cache software on remote computers that are caching data from the given I/O device.” IBM proposes the corresponding structure should be “the ‘write invalidate,’ ‘message receive,’ and ‘cache data invalidate’ program flows, steps and data structures disclosed or referred to by Figs. 5L, 5M, 5N, and 5P, and col. 22, 1.23-col. 24, l. 9 of the ‘226 patent.”

SuperSpeed argues that IBM’s proposed structure is too broad in that the subroutines it references include structure that is not necessary for the agreed function of the term. SuperSpeed argues, for instance, the “cache data invalidate” subroutine includes structure for hashing disk blocks, selecting hash chains, removing previously invalidated TCMBs from the hash chain and the least recently used queue, and incrementing the counter keeping track of how many times data had been invalidated. ‘226 22:32-41, 22:56-65; Fig 5L, boxes 592-594, 604-606 & 610. Further, SuperSpeed argues, the “write invalidate” subroutine has structure for intercepting write completion instructions from an I/O device. ‘226 23:32-35; Fig. 5M, box 628. None of these structures, SuperSpeed argues, relates to *using computer communication channels*, much less using computer communication channels to invalidate data in remote caches. Bennett Dec., ¶ 64.

IBM's construction is too broad, in that it not only includes structure for *using the communication channels* but also to perform the actual invalidation. That structure is unnecessary. As construed by the court, "invalidate" requires an indication of data modification. The structure, therefore, is limited to what is needed to facilitate the use of the communication channels to send the indication.

SuperSpeed's proposal is sufficiently descriptive under *Harris*, 417 F.3d at 1254. The corresponding structure, therefore, is "a computer executing a software algorithm that sends a message indicating the invalidity of certain data found on an I/O device to the cache software on remote computers that are caching data from the given I/O device and its equivalents."

12. "means for listening on said network for a request from a new computer to connect to said network"

Both parties agree this is a mean-plus-function claim limitation and that the recited function should be given the meaning "listening on said network for a request from a new computer to connect to said network." The parties disagree, however, on the corresponding structure. SuperSpeed proposes the corresponding structure should be "a computer executing software that configures the operating system to execute code in the cache driver in response to a connection request from another computer." IBM proposes that there is no disclosed corresponding structure, rendering the claim indefinite.

It is well settled that the patent's specification must disclose some structure to perform the recited function. SuperSpeed is unable to identify any structure in the specification that corresponds to the recited function. *Default Proof Credit Card Sys., Inc. v. Home Depot U.S.A., Inc.*, 412 F.3d 1291, 1302 (Fed.Cir.2005). SuperSpeed is unable to point to any details or algorithms for listening

on the network. *See WMS Gaming, Inc. v. Int'l Game Tech.*, 184 F.3d 1339, 1349 (Fed. Cir. 1999) (“In a means-plus-function claim in which the disclosed structure is a computer, or microprocessor, programmed to carry out an algorithm, the disclosed structure is not the general purpose computer, but rather the special purpose computer programmed to perform the disclosed algorithm”). The court, therefore, holds that claims 1, 16, 27, and 37 of the ‘244 patent are indefinite. *See Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1367 (Fed. Cir. 2008) (affirming district court’s finding of indefiniteness where the claim recited a “bank computer” but nothing in the written description expressly described what was going on inside that bank computer).

V. Conclusion

The court adopts the constructions set forth in this opinion for the disputed terms of the ‘226, ‘244, ‘013, ‘767, and ‘129 patents. The parties are ordered that they may not refer, directly or indirectly, to each other’s claim construction positions in the presence of the jury. Likewise, the parties are ordered to refrain from mentioning any portion of this opinion, other than the actual definitions adopted by the Court, in the presence of the jury. Any reference to claim construction proceedings is limited to informing the jury of the definitions adopted by the Court.

SIGNED this 11th day of February, 2009.



CHARLES EVERINGHAM IV
UNITED STATES MAGISTRATE JUDGE